arXiv:1610.08713v1 [cs.SE] 27 Oct 2016

The Probabilistic Model Checker Storm
(Extended Abstract)

Christian Dehnert, Sebastian Junges, Joost-Pieter Katoen, and Matthias Volk

RWTH Aachen University

We present a new probabilistic model checker Storm. Using state-of-the-art libraries, we aim for both high performance and versatility. This extended abstract gives a brief overview of the features of Storm.

Goals. Probabilistic model checking has matured immensely in the past decade [1,2,3,4]. Applications and uses go far beyond the standard algorithms and objectives. We therefore developed an extensible toolbox for probabilistic model checking that offers reusable modules to quickly implement new functionality. The rise of more and more specialized tools witnesses that a single tool is rarely suitable for the broad range of applications and models. Consequently, we offer a range of engines that make use of various backend solvers. This core is aiming for high performance and low memory requirements. On top of it we realized sophisticated techniques ranging from generation of probabilistic counterexamples [5], over permissive scheduler generation [6] and parameter synthesis [7,8] to the analysis of dynamic fault trees [9] and probabilistic program verification [10]. Finally, the wide range of input languages supports the development of new techniques in a variety of fields.

Supported model types. Storm supports both discrete-time and continuous-time Markov models and nondeterministic variants thereof. Table 1 gives an overview and classification about the concrete model types: discrete-time Markov chains (DTMCs), continuous-time Markov chains (CTMCs), Markov decision processes (MDPs) and Markov automata (MA), the compositional variant of continuous-time Markov Decision processes (CTMDPs) [4]. Additionally, all these model types can be enriched with reward models.

Table 1. Model types supported by Storm.

	discrete-time	continuous-time
deterministic	DTMCs	CTMCs
nondeterministic	MDPs	MA

Modelling languages. There are several modelling languages that can be used to specify the aforementioned model types. First of all, in the spirit of MRMC [11], models can be passed in a format that explicitly enumerates transitions. However, STORM also supports a variety of symbolic input formats. Most prominently, the PRISM input language [12] is supported, which enables us to consider all benchmark models from the PRISM benchmark suite [13] that represent supported model types. In an attempt to unify the probabilistic modelling language landscape, the JANI format [14] is currently under development and its first version is included in STORM. Besides, it is the first tool to support *every* generalized stochastic Petri net (GSPN) [15] via both a dedicated model builder as well as an encoding in JANI. Additionally, STORM features the analysis of dynamic fault trees (DFTs) [16] and has been shown to outperform competing tools in this domain [9]. Finally, we support probabilistic programs written in the conditional probabilistic guarded command language (cpGCL) [17], again via a reduction to JANI.

Properties. The main focus of our tool is probabilistic branching time logics, i.e. PCTL [18] and CSL [19] for discrete-time and continuous-time models, respectively. To enable the treatment of reward-objectives, we support reward extensions of these logics in a similar way as PRISM. As conditional probabilities and conditional rewards [20] have proven to express interesting properties, we also support these.

Engines. There are different representations of probabilistic models in memory that differ in efficiency depending on the characteristics of the input model. STORM features two distinct representations: while *sparse matrices* tend to work well for small and moderately sized models, multi-terminal decision diagrams (MTBDDs) are able to represent gigantic systems. To enable the treatment of a broader class of input models, the user can select between several engines built around the two in-memory representations to perform the verification tasks. Both STORM's *sparse* and the *learning* engine purely use a sparse matrix based representation. While the former amounts to an efficient implementation of the standard approaches, the latter one implements the ideas of [21] to provide sound statistical model checking for discrete-time models. Two other engines, *dd* and *hybrid*, use MTBDDs as the primary representation. While *dd* exclusively uses decision diagrams, *hybrid* also uses sparse matrices for operations deemed more suitable on this data format.

Solvers. STORM's infrastructure is built around the notion of a *solver*. For instance, solvers are available for sets of linear or Bellman equations (both using sparse matrices as well as MTBDDs), (mixed-integer) linear programming (MILP) and satisfiability modulo theories (SMT) solving. Offering these interfaces has several key advantages. First, it provides easy and coherent access to the tasks commonly involved in probabilistic model checking. Secondly, it enables the use of dedicated state-of-the-art high-performance libraries for the task at hand. More specifically, as the performance characteristics of different

Table 2. Solvers offered by STORM.

solver type	available solvers
linear equations (sparse)	Eigen [22], gmm++ [23], built-in
linear equations (MTBDD)	CUDD [24], Sylvan [25]
Bellman equations (sparse)	Eigen, gmm++, built-in
Bellman equations (MTBDD)	CUDD, Sylvan
(MI)LP	Gurobi [26], glpk [27]
SMT	Z3 [28], MathSAT [29], SMTLIB [30]

backend solvers can vary drastically for the same input, this permits choosing the best solver for a given task. Licensing problems are avoided, because implementations can be easily enabled and disabled, depending on whether or not the particular license fits the requirements. Finally, implementing new solver functionality is easy and can be done without knowledge about the global code base. For each of those interfaces, several actual implementations exist. Table 2 gives an overview over the currently available implementations. Many components in STORM make use of these solvers. The most prominent example is the obvious use of the equation solvers for answering standard verification queries. However, various other modules use them too, for example model generation (SMT), counterexample generation [31,5] (SMT, MILP) and permissive scheduler generation [32,6] (SMT, MILP).

Parametric models and exact arithmetic. STORM was used as backend in [7,8]. By using the dedicated library CARL [33] for the representation of rational functions and applying novel algorithms for the analysis of parametric discrete-time models, it has proven to significantly outperform other tools. In addition, for these models STORM is able to compute exact solutions rather than using floating point arithmetic.

Usage. STORM can be used via three interfaces. For end-users we provide command-line interfaces: There are several binaries that provide specialized access to the available settings for different tasks. For example, STORM-DFT offers the features and settings related to the analysis of DFTs. Advanced users can utilize one of the many settings to tune the performance. Developers may either use a C++ API that offers fine-grained and performance-oriented access to STORM's functionality, or use a Python API which supports rapid prototyping and allows users to profit from high-performance implementations within STORM.

Related work. There are several tools whose functionality overlaps with that of STORM (in alphabetical order): FACT [34], FIG [35], IMCA [36], ISCASMC [37], MODEST [38], MRMC [11], PARAM [39], PASS [40], PAT [41], PRISM

[42] and UPPAAL SMC [43]. Moreover, there are specialized tools for the generation of probabilistic counterexamples (COMICS [44]), DFTs (DFTCALC [45]) and GSPNs (GREATSPN [46]). Despite the existence of these tools, we believe that STORM is unique in its trade-off between performance and modularity, the supported solvers and wide range of supported modelling languages.

References

1. Kwiatkowska, M., Norman, G., Parker, D.: Stochastic model checking. Proc. of SFM. Vol. 4486 of LNCS, Springer (2007) 220–270
2. Kwiatkowska, M.: Model checking for probability and time: from theory to practice. LICS, IEEE CS (2003) 351
3. Baier, C.: Probabilistic model checking. Dependable Software Systems Engineering. Vol. 45 of NATO Science for Peace and Security Series - D: Information and Communication Security. IOS Press (2016) 1–23
4. Katoen, J.P.: The probabilistic model checking landscape. Proc. of LICS, ACM (2016) 31–46
5. Dehnert, C., Jansen, N., Wimmer, R., Ábrahám, E., Katoen, J.P.: Fast debugging of PRISM models. Proc. of ATVA. Vol. 8837 of LNCS, Springer (2014) 146–162
6. Junges, S., Jansen, N., Dehnert, C., Topcu, U., Katoen, J.P.: Safety-constrained reinforcement learning for MDPs. Proc. of TACAS. (2016)
7. Dehnert, C., Junges, S., Jansen, N., Corzilius, F., Volk, M., Bruintjes, H., Katoen, J.P., Ábrahám, E.: Prophesy: A probabilistic parameter synthesis tool. Proc. of CAV. Vol. 9206 of LNCS, Springer (2015) 214–231
8. Quatmann, T., Dehnert, C., Jansen, N., Junges, S., Katoen, J.P.: Parameter synthesis for Markov models: Faster than ever. Proc. of ATVA. Vol. 9938 of LNCS (2016) 50–67
9. Volk, M., Junges, S., Katoen, J.P.: Advancing dynamic fault tree analysis - get succinct state spaces fast and synthesise failure rates. Proc. of SAFECOMP. Vol. 9922 of LNCS, Springer (2016) 253–265
10. Jansen, N., Dehnert, C., Kaminski, B.L., Katoen, J.P., Westhofen, L.: Bounded model checking for probabilistic programs. Proc. of ATVA. Vol. 9938 of LNCS (2016) 68–85
11. Katoen, J.P., Zapreev, I.S., Hahn, E.M., Hermanns, H., Jansen, D.N.: The ins and outs of the probabilistic model checker MRMC. Performance Evaluation **68**(2) (2011) 90–104
12. PRISM Website: (2015) `http://prismmodelchecker.org`.
13. Kwiatkowska, M., Norman, G., Parker, D.: The PRISM benchmark suite. Proc. of QEST, IEEE CS (2012) 203–204
14. Budde, C.E., Dehnert, C., Hahn, E.M., Hartmanns, A., Junges, S., Turrini, A.: Jani: Quantitative model and tool interaction. (in preparation)
15. Eisentraut, C., Hermanns, H., Katoen, J.P., Zhang, L.: A semantics for every GSPN. Proc. of Petri Nets. Vol. 7927 of LNCS, Springer (2013) 90–109
16. Stamatelatos, M., Vesely, W., Dugan, J.B., Fragola, J., Minarick, J., Railsback, J.: Fault Tree Handbook with Aerospace Applications. NASA Headquarters (2002)
17. McIver, A., Morgan, C.: Abstraction, Refinement and Proof for Probabilistic Systems. Monographs in Computer Science. Springer (2005)
18. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. Formal Aspects of Computing **6**(5) (1994) 512–535
19. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.K.: Verifying continuous time Markov chains. Proc. of CAV. Vol. 1102 of LNCS, Springer (1996) 269–276
20. Baier, C., Klein, J., Klüppelholz, S., Märcker, S.: Computing conditional probabilities in Markovian models efficiently. Proc. of TACAS. Vol. 8413 of LNCS, Springer (2014) 515–530
21. Brázdil, T., Chatterjee, K., Chmelik, M., Forejt, V., Kretínský, J., Kwiatkowska, M., Parker, D., Ujma, M.: Verification of Markov decision processes using learning algorithms. Proc. of ATVA. Vol. 8837 of LNCS, Springer (2014) 98–114

22. Guennebaud, G., Jacob, B., et al.: Eigen v3. (http://eigen.tuxfamily.org)
23. GMM++ Website: (2015) http://download.gna.org/getfem/html/homepage/gmm/index.html.
24. CUDD Website: (2015) http://vlsi.colorado.edu/~fabio/CUDD/html/index.html.
25. van Dijk, T., van de Pol, J.: Sylvan: Multi-core decision diagrams. Proc. of TACAS. Vol. 9035 of LNCS, Springer (2015) 677–691
26. Gurobi Optimization, Inc.: Gurobi optimizer reference manual. http://www.gurobi.com (2015)
27. GNU project: Linear programming kit, version 4.6. http://www.gnu.org/software/glpk/glpk.html (2016)
28. de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. Proc. of TACAS. Vol. 4963 of LNCS, Springer (2008) 337–340
29. Cimatti, A., Griggio, A., Schaafsma, B.J., Sebastiani, R.: The mathsat5 SMT solver. Proc. of TACAS. Vol. 7795 of Lecture Notes in Computer Science, Springer (2013) 93–107
30. Barrett, C., Stump, A., Tinelli, C., Boehme, S., Cok, D., Deharbe, D., Dutertre, B., Fontaine, P., Ganesh, V., Griggio, A., Grundy, J., Jackson, P., Oliveras, A., KrstiÄĞ, S., Moskal, M., Moura, L.D., Sebastiani, R., Cok, T.D., Hoenicke, J.: C.: The smt-lib standard: Version 2.0. Technical report (2010)
31. Wimmer, R., Jansen, N., Vorpahl, A., Ábrahám, E., Katoen, J.P., Becker, B.: High-level counterexamples for probabilistic automata. Logical Methods in Computer Science **11**(1) (2015)
32. Dräger, K., Forejt, V., Kwiatkowska, M., Parker, D., Ujma, M.: Permissive controller synthesis for probabilistic systems. Logical Methods in Computer Science **11**(2) (2015)
33. CArL Website: (2015) http://goo.gl/8QsVxv.
34. Calinescu, R., Johnson, K., Paterson, C.: FACT: A probabilistic model checker for formal verification with confidence intervals. Proc. of TACAS. Vol. 9636 of LNCS, Springer (2016) 540–546
35. Budde, C.E., D'Argenio, P.R., Monti, R.E.: Compositional construction of importance functions in fully automated importance splitting. Proc. of VALUETOOLS, ICST (2016)
36. Guck, D., Timmer, M., Hatefi, H., Ruijters, E., Stoelinga, M.: Modelling and analysis of Markov reward automata. Proc. of ATVA. Vol. 8837 of LNCS, Springer (2014) 168–184
37. Hahn, E.M., Li, Y., Schewe, S., Turrini, A., Zhang, L.: iscasMc: A web-based probabilistic model checker. Proc. of FM. Vol. 8442 of LNCS, Springer (2014) 312–317
38. Hartmanns, A., Hermanns, H.: The modest toolset: An integrated environment for quantitative modelling and verification. Proc. of TACAS. Vol. 8413 of LNCS, Springer (2014) 593–598
39. Hahn, E.M., Hermanns, H., Wachter, B., Zhang, L.: PARAM: A model checker for parametric Markov models. Proc. of CAV. Vol. 6174 of LNCS, Springer (2010) 660–664
40. Hahn, E.M., Hermanns, H., Wachter, B., Zhang, L.: PASS: abstraction refinement for infinite probabilistic models. Proc. of TACAS. Vol. 6015 of LNCS, Springer (2010) 353–357
41. Liu, Y., Sun, J., Dong, J.S.: PAT 3: An extensible architecture for building multi-domain model checkers. Proc. of ISSRE, IEEE CS (2011) 190–199
42. Kwiatkowska, M., Norman, G., Parker, D.: Prism 4.0: Verification of probabilistic real-time systems. Proc. of CAV. Vol. 6806 of LNCS, Springer (2011) 585–591

43. Bulychev, P.E., David, A., Larsen, K.G., Mikucionis, M., Poulsen, D.B., Legay, A., Wang, Z.: UPPAAL-SMC: statistical model checking for priced timed automata. Proc. of QAPL. Vol. 85 of EPTCS (2012) 1–16
44. Jansen, N., Ábrahám, E., Volk, M., Wimmer, R., Katoen, J.P., Becker, B.: The COMICS tool – Computing minimal counterexamples for DTMCs. Proc. of ATVA. Vol. 7561 of LNCS, Springer (2012) 349–353 (to appear).
45. Arnold, F., Belinfante, A., van der Berg, F., Guck, D., Stoelinga, M.: DFTCalc: A tool for efficient fault tree analysis. Proc. of SAFECOMP. Vol. 8153 of LNCS, Springer (2013) 293–301
46. Amparore, E.G., Beccuti, M., Donatelli, S.: (stochastic) model checking in GreatSPN. Proc. of Petri Nets. Vol. 8489 of LNCS, Springer (2014) 354–363